

Integrated circuit and method for sending requests

The invention relates to an integrated circuit as described in the introductory part of claim 1. The invention also relates to a method for sending requests as described in the introductory part of claim 7.

5

A request-response transaction model is frequently used as a communication model for systems on integrated circuits. The transaction model can either be deployed in systems using a bus architecture or in systems using a network architecture, in order to establish communication between the modules. Using such a transaction model in a network on an integrated circuit provides backward compatibility with existing interconnects, for example buses.

The transaction model uses requests and responses. A request comprises a command (e.g. read, write) with parameters, such as an address or a burst length, and optionally the request comprises a data part. Responses carry an acknowledgement indicating the result of the execution of a request, and optionally they carry a data part.

Another communication model is the message-passing model, which uses messages and acknowledgements. Such an acknowledgement indicates the receipt of a message rather than the execution of a request.

In networks on an integrated circuit a first module (also referred to as master, master module or initiator) typically has access to an address space, wherein addresses identify locations within second modules (also referred to as slaves, slave modules or targets). Depending on circumstances, it may be necessary to address two or more second modules simultaneously. For example, this is needed when the execution of a request by two or more second modules needs to be started simultaneously and the starting is performed by writing to start registers mapped in the address space; all these start registers should be written to at the same time. Another example is when data is replicated to different memories to be processed locally. In these cases the first module replicates the request and the resulting plurality of replicated requests is sent to the second modules. This has the disadvantage that the first module cannot send a request to more than one second module

using a single address, but it must replicate the request and send the replicated requests to the second modules using a different address for each second module. This causes a large burden on the first module.

5

It is an object of the invention to provide an integrated circuit and a method of the kind set forth which reduce the burden on the first module. In order to achieve the said object the integrated circuit is characterized by the characterizing portion of claim 1 and the method is characterized by the characterizing portion of claim 7.

10

The burden on the first module can be reduced by providing a network which is capable of replicating a request into at least two replicated requests, and which is capable of sending the replicated requests to the second modules. If the network can perform these tasks, then the first module can be relieved of them.

An embodiment is defined in claim 2, wherein the network comprises a facility for mapping at least one special address (also referred to as multicast address) onto at least two further addresses. This enables the first module to send a single request to a single address instead of replicating the request and sending the replicated requests to various addresses.

It is also possible to map one or more multicast addresses onto one or more other multicast addresses; this embodiment is defined in claim 3. This has the constraint that no recurrence should occur.

Depending on circumstances, it is convenient to specify a range of multicast addresses once instead of specifying a number of separate multicast addresses. The embodiment defined in claim 4 provides a facility for defining such a range of multicast addresses.

Another embodiment is defined in claim 5, wherein a multicast connection is deployed to relieve the first module of the replication and dispatch tasks. The first module can send a single request comprising a connection identifier referring to such a connection; the network then replicates the single request into at least two replicated requests and sends the replicated requests through the connection to the second modules.

One or more dedicated nodes in the network may be used to replicate the single request and send the replicated requests. The embodiment defined in claim 6 comprises a network interface to replicate the single request and send the replicated requests.

The invention overcomes the shortcomings of multicast transactions in networks on an integrated circuit, because the network can provide a multicast request to at least two second modules in response to a single request from the first module.

It is noted that US 2002/0093964 discloses a protocol for routers (data
5 switching nodes) and supervisors to exchange data. The router can send commands to the supervisor including a learn/delete/search multicast address command. The supervisor provides information to the router about multicast packets that must be routed. However, the supervisor does not perform an actual multicast; the router must perform this multicast. A method of multicasting of the kind set forth is not disclosed in US 2002/0093964.

10

The present invention is described in more detail with reference to the drawings, in which:

Fig. 1 illustrates a network on an integrated circuit;

15

Fig. 2 illustrates a method of multicasting wherein at least two replicated requests are sent from a first module to at least two second modules;

Fig. 3 illustrates how addresses are used to address the second modules;

Fig. 4 illustrates a method of multicasting according to the invention;

Fig. 5 illustrates how multicast addresses are used according to the invention;

20

Fig. 6 illustrates a multicast range;

Fig. 7 illustrates a multicast connection according to the invention.

Fig. 1 schematically shows an integrated circuit IC which deploys a network
25 for communication between a plurality of modules M_1 , M_2 , M_3 up to and including M_n . Examples of modules are central processing units (CPUs), application specific processors, memories and memory controllers. The network comprises nodes N_1 , N_2 up to and including N_x , and connections between the nodes. This network architecture provides the interconnect between the modules and can be deployed as an alternative for the conventional bus
30 architecture on an integrated circuit.

Fig. 2 illustrates a method of multicasting wherein at least two replicated requests REQ_2 , REQ_3 up to and including REQ_n are sent from a first module M_1 to at least two second modules M_2 , M_3 up to and including M_n . The second modules M_2 , M_3 up to and including M_n send back responses $RESP_2$, $RESP_3$ up to and including $RESP_n$ to the first

module M_1 . The disadvantage of this method of multicasting is that the first module M_1 cannot send a request to the two second modules M_2 , M_3 up to and including M_n using a single address, but it must replicate the request and send the replicated requests REQ_2 , REQ_3 up to and including REQ_n to the second modules, using a different address for each second module. This leads to a large burden on the first module M_1 .

Fig. 3 illustrates how addresses can be used to address the second modules M_2 , M_3 up to and including M_n . An address space ADDR_SPC comprises a range of addresses 0 up to and including EFFF. A first sub range thereof, comprising addresses 0 up to and including AFFF, is associated with second module M_2 , and a second sub range, comprising addresses B000 up to and including EFFF, is associated with second module M_3 . In order to send a request to both second modules M_2 and M_3 , the first module M_1 must replicate the request and send one replicated request to an address in the first sub range 0 up to and including AFFF, for example to address 3A98, and another replicated request to an address in the second sub range B000 up to and including EFFF, for example to address C350.

Fig. 4 illustrates a method of multicasting according to the invention. Instead of sending replicated requests to the second modules M_2 , M_3 up to and including M_n , the first module M_1 can send a single request SREQ which is replicated by the network, and distributed to two or more second modules. This can be realized, for example, using one or more special addresses to specify the addresses of the second modules. Such a special address, also referred to as a multicast address, is mapped onto the addresses of the second modules M_2 , M_3 up to and including M_n in the address space. The network performs the replication of the single request SREQ, resulting in at least two replicated requests $SREQ_2$, $SREQ_3$ up to and including $SREQ_n$, which are sent to the second modules M_2 , M_3 up to and including M_n . The network may deploy a network interface for this purpose; alternatively one or more dedicated nodes may be used.

Fig. 5 illustrates a mapping between a multicast address and two addresses which are associated with the second modules M_2 and M_3 respectively. The multicast address must be configured such that it is associated with at least one of the addresses of each of the second modules M_2 and M_3 ; this may be done by the boot code which configures the network or it may be done at runtime. In this example the address space ADDR_SPC comprises a range of regular addresses 0 up to and including EFFF, which are associated directly with the second modules. It is assumed that a network interface NI, which is comprised in the network, performs the replication of a request and sends the replicated requests to the second modules M_2 and M_3 .

A special address F000, also referred to as a multicast address, is mapped onto two regular addresses: address 3A98 which is in the sub range 0 up to and including AFFF associated with second module M_2 , and address C350 which is in the sub range B000 up to and including EFFF associated with second module M_3 . Such a mapping may take place by using a lookup table or by a logical operation, for example. Now the first module M_1 can send a request to the multicast address F000, and then the request is replicated by the network interface NI and sent to the addresses 3A98 and C350, which are associated with second module M_2 and second module M_3 respectively.

Note that it is possible that a multicast address is mapped onto another multicast address, but recurrence must be avoided in the sense that a first multicast address should not be mapped onto a second multicast address which in turn is mapped onto the first multicast address. Note also that it is possible that a multicast address is mapped onto two or more addresses within a single second module.

In addition, it is possible to define a multicast range, which is illustrated in Fig. 6. Instead of specifying a large number of multicast addresses which are mapped onto regular addresses corresponding to second modules M_2 and M_3 , it is possible to specify once a complete range of multicast addresses. In this example the range of multicast addresses $F000+x$, where 'x' is a variable having values in the range 0 up to and including 3FF, is mapped onto addresses $1000+x$ (which correspond to second module M_2) and onto addresses $3000+x$ (which correspond to second module M_3). Hence, multicast address F000 is mapped onto addresses 1000 and 3000, multicast address F001 is mapped onto addresses 1001 and 3001, multicast address F002 is mapped onto addresses 1002 and 3002, etc. The last multicast address in the range is F3FF, which is mapped onto addresses 13FF and 33FF. This multicast range has the advantage that it is not necessary to specify 1024 separate multicast addresses.

Alternatively, a multicast connection can be deployed to reduce the burden on the first module M_1 , which is illustrated in Fig. 7. In general, connections are used in networks to describe and identify communication with different properties, such as guaranteed throughput, latency and jitter, ordered delivery, or flow control. In this context, a connection is used to identify a first module and a second module or a number of second modules. The connection comprises physical means and control information required to enable a transaction between the first module and the second module(s). The path between the first module and the second module(s) can be determined at runtime, at (re)configuration time and/or it can be predetermined by the boot code. The required control information

comprises a mapping between a connection identifier and one or more network interface ports (NIPs). The control information may be stored in the network nodes (i.e. routers and network interfaces) and/or it may be contained in the header of a packet, for example.

In this example a multicast connection is set up from a first module M_1 to two
5 second modules M_2 and M_3 . Connections require that a connection identifier CID is sent along with a request. A request on such a connection can then automatically be sent to all the second modules of the connection. Note that the requests still carry addresses which are used as internal addresses for the second modules, i.e. addresses which identify locations within the second modules, but which are not used to replicate and distribute requests to the second
10 modules.

A connection is set up during a configuration stage of the network. Typically the network is configured by the boot code but it may also be configured at runtime. The connection identifier CID has a value which identifies the connection; in this case the value is '0'. In the example shown the value of the connection identifier CID is mapped onto network
15 interface ports NIP_2 and NIP_3 via the mapping $0 \rightarrow \{NIP_2, NIP_3\}$. The network interface ports NIP_2 and NIP_3 form part of network interfaces NI_2 and NI_3 respectively; note that one network interface may have more than one network interface port and several network interface ports may be associated with a single address. These network interface ports NIP_2 and NIP_3 are in turn associated with the addresses of the second modules M_2 and M_3 . Now
20 the connection identifier CID can be sent along with a request; the network interface NI_1 can replicate the request and send the replicated requests through the connection. Via the network interface ports NIP_2 and NIP_3 the request can be delivered at the addresses of the second modules M_2 and M_3 .

It is remarked that the scope of protection of the invention is not restricted to
25 the embodiments described herein. Neither is the scope of protection of the invention restricted by the reference symbols in the claims. The word 'comprising' does not exclude other parts than those mentioned in a claim. The word 'a(n)' preceding an element does not exclude a plurality of those elements. Means forming part of the invention may both be implemented in the form of dedicated hardware or in the form of a programmed general-
30 purpose processor. The invention resides in each new feature or combination of features.